

A new integral loss function for Bayesian optimization

Emmanuel Vazquez and Julien Bect

SUPELEC, 3 rue Joliot-Curie, 91192 Gif-sur-Yvette, France
email: {firstname}.{lastname}@supelec.fr

Abstract

We consider the problem of maximizing a real-valued continuous function f using a Bayesian approach. Since the early work of Jonas Mockus and Antanas Žilinskas in the 70's, the problem of optimization is usually formulated by considering the loss function $\max f - M_n$ (where M_n denotes the best function value observed after n evaluations of f). This loss function puts emphasis on the value of the maximum, at the expense of the location of the maximizer. In the special case of a one-step Bayes-optimal strategy, it leads to the classical Expected Improvement (EI) sampling criterion. This is a special case of a Stepwise Uncertainty Reduction (SUR) strategy, where the risk associated to a certain uncertainty measure (here, the expected loss) on the quantity of interest is minimized at each step of the algorithm. In this article, assuming that f is defined over a measure space (\mathbb{X}, λ) , we propose to consider instead the integral loss function $\int_{\mathbb{X}} (f - M_n)_+ d\lambda$, and we show that this leads, in the case of a Gaussian process prior, to a new numerically tractable sampling criterion that we call EI² (for Expected Integrated Expected Improvement). A numerical experiment illustrates that a SUR strategy based on this new sampling criterion reduces the error on both the value and the location of the maximizer faster than the EI-based strategy.

Keywords: Bayesian optimization, computer experiments, Gaussian process, global optimization, sequential design
62L05; 62M20; 62K20; 60G15; 60G25; 90C99

1. Introduction

Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be a real-valued continuous function defined on a compact subset \mathbb{X} of \mathbb{R}^d , $d \geq 1$. We consider the problem of finding an approximation of the maximum of f ,

$$M = \max_{x \in \mathbb{X}} f(x),$$

and of the set of maximizers,

$$x^* \in \operatorname{argmax}_{x \in \mathbb{X}} f(x),$$

using a sequence of queries of the value of f at points $X_1, X_2, \dots \in \mathbb{X}$. At iteration $n + 1$, the choice of the evaluation point X_{n+1} is allowed to depend on the results $f(X_1), \dots, f(X_n)$ of the evaluation of f at X_1, \dots, X_n . Thus, the construction of an optimization strategy $\underline{X} = (X_1, X_2, \dots)$ can be seen as a sequential decision problem.

We adopt the following Bayesian approach for constructing \underline{X} . The unknown function f is considered as a sample path of a random process ξ defined on some probability space $(\Omega, \mathcal{B}, \mathbf{P}_0)$, with parameter $x \in \mathbb{X}$. For a given f , the efficiency of a strategy \underline{X} can be measured in different ways. For instance, a natural loss function for measuring the performance of \underline{X} at iteration n is

$$\varepsilon_n(\underline{X}, f) = M - M_n, \quad (1)$$

with $M_n = \max(f(X_1), \dots, f(X_n))$. The choice of a loss function ε_n , together with a random process model, makes it possible to define the following one-step Bayes-optimal strategy:

$$\begin{cases} X_1 = x_{\text{init}} \\ X_{n+1} = \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathbf{E}_n(\varepsilon_{n+1}(\underline{X}, \xi) | X_{n+1} = x_{n+1}), \quad \forall n \geq 1, \end{cases} \quad (2)$$

where \mathbb{E}_n denotes the conditional expectation with respect to the σ -algebra \mathcal{F}_n generated by the random variables $X_1, \xi(X_1), \dots, X_n, \xi(X_n)$. This Bayesian decision-theoretic point of view has been initiated during the 70's by the work of Jonas Mockus and Antanas Žilinskas (see Mockus et al., 1978; Mockus, 1989, and references therein).

For instance, consider the loss defined by (1). Then, at iteration $n + 1$, the strategy (2) can be written as

$$\begin{aligned} X_{n+1} &= \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathbb{E}_n(M - M_{n+1} \mid X_{n+1} = x_{n+1}) \\ &= \operatorname{argmax}_{x_{n+1} \in \mathbb{X}} \rho_n(x_{n+1}), \end{aligned} \quad (3)$$

where $\rho_n(x) := \mathbb{E}_n(\max(\xi(x) - M_n, 0))$ is the *Expected Improvement* (EI) criterion, introduced by Mockus et al. (1978) and later popularized through the EGO algorithm (Jones et al., 1998), both in the case of Gaussian process models (for which $\rho_n(x)$ admits a closed-form expression as a function of the posterior mean and variance of ξ at x).

The contribution of this paper is a new loss function for evaluating the efficiency of an optimization strategy, from which we can derive, in the case of a Gaussian process prior, a numerically tractable sampling criterion for choosing the evaluations points according to a one-step Bayes-optimal strategy. Section 2 explains our motivation for the introduction of a novel loss function, and then proceeds to present the loss function itself and the associated sampling criterion. The numerical implementation of this new sampling criterion is discussed in Section 3. Finally, Section 4 presents a one-dimensional example that illustrates qualitatively the effect of using our new loss function, together with a numerical study that assesses the performance of the criterion from a statistical point of view on a set of sample paths of a Gaussian process.

2. An integral loss function

Observe that (3) can be rewritten as

$$X_{n+1} = \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathbb{E}_n(H_{n+1} \mid X_{n+1} = x_{n+1}), \quad (4)$$

with $H_n = \mathbb{E}_n(M - M_n)$. The \mathcal{F}_{n+1} -measurable random variable H_{n+1} in the right-hand side of (4) can be seen as a measure of the uncertainty about M at iteration $n+1$: indeed, according to Markov's inequality, $M \in [M_{n+1}; M_{n+1} + H_{n+1}/\delta]$ with probability at least $1 - \delta$ under \mathbb{P}_{n+1} . Thus, this strategy is actually a special case of *stepwise uncertainty reduction* (Villemonteix et al., 2009; Bect et al., 2012; Chevalier et al., 2013).

In a global optimization problem, it is generally of interest to obtain a good approximation of *both* M and x^* . The classical loss function $\varepsilon_n = M - M_n$ is not very satisfactory from this respect, since the associated uncertainty measure $H_n = \mathbb{E}_n(M - M_n)$ puts all the emphasis on M , at the expense of x^* . Other uncertainty measures have been proposed recently, which take the opposite approach and focus on x^* only (Villemonteix et al., 2009; Picheny, 2014a,b).

Assume now that \mathbb{X} is endowed with a finite positive measure λ (e.g., Lebesgue's measure restricted to \mathbb{X}), and let us remark that the classical loss function (1) is proportional to $\lambda(\mathbb{X})(M - M_n)$, that is, to the area of the hatched region in Figure 1a. This illustrates that $H_n = \mathbb{E}_n(\varepsilon_n)$ is only a coarse measure of the uncertainty about the pair (M, x^*) . We propose to use instead the integral loss function

$$\varepsilon'_n(\underline{X}, f) = \int_{\mathbb{X}} (f(x) - M_n)_+ \lambda(dx), \quad (5)$$

where $z_+ := \max(z, 0)$. This new loss function is depicted in Figure 1b. The associated uncertainty measure $H'_n = \mathbb{E}_n(\varepsilon'_n)$ should, intuitively, provide a finer measure of the uncertainty about the pair (M, x^*) and thereby lead to better optimization algorithms. The corresponding stepwise uncertainty reduction strategy can be written as

$$\begin{aligned} X_{n+1} &= \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathbb{E}_n \left(\int_{\mathbb{X}} (\xi(y) - M_{n+1})_+ \lambda(dy) \mid X_{n+1} = x_{n+1} \right) \\ &= \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathbb{E}_n \left(\int_{\mathbb{X}} \mathbb{E}_{n+1}((\xi(y) - M_{n+1})_+) \lambda(dy) \mid X_{n+1} = x_{n+1} \right) \\ &= \operatorname{argmin}_{x_{n+1} \in \mathbb{X}} \mathfrak{F}_n(x_{n+1}), \end{aligned} \quad (6)$$

where

$$\mathfrak{S}_n(x_{n+1}) := \mathbb{E}_n \left(\int_{\mathbb{X}} \rho_{n+1}(y) \lambda(dy) \mid X_{n+1} = x_{n+1} \right) \quad (7)$$

is a new sampling criterion than we call EI^2 (for Expected Integrated Expected Improvement). Note that the strategy (6) is very different in spirit from the classical one, associated to the EI criterion. Indeed, while the classical strategy selects a point where the *current* EI is *maximal*, the new strategy selects a point where the integral of the *future* EI is *minimal*, in expectation.

Remark. The sampling criterion defined by (7) is a one-point sampling criterion; that is, a sampling criterion for use in a fully sequential setting. A multi-point sampling criterion can be defined similarly, for use in a batch-sequential setting:

$$\mathfrak{S}_{n,r}(x_{n+1}, \dots, x_{n+r}) := \mathbb{E}_n \left(\int_{\mathbb{X}} \rho_{n+r}(y) \lambda(dy) \mid X_{n+1} = x_{n+1}, \dots, X_{n+r} = x_{n+r} \right) \quad (8)$$

(see Chevalier and Ginsbourger (2013); Chevalier et al. (2013) and references therein for more information on multi-point stepwise uncertainty reduction strategies).




Figure 1: A diagrammatic interpretation of the loss functions ε_n (left plot) and ε'_n (right plot).

3. Numerical approximation of the sampling criterion

Numerical approximations of the sampling criterion \mathfrak{S}_n can be obtained with an acceptable computational complexity when ξ is a Gaussian process. Rewrite (7) as

$$\mathfrak{S}_n(x_{n+1}) = \int_{\mathbb{X}} \bar{\rho}_n(y; x_{n+1}) \lambda(dy), \quad (9)$$

where $\bar{\rho}_n(y; x_{n+1})$, which we shall call the *Expected Expected Improvement* (EEI) at $y \in \mathbb{X}$ given a new evaluation at $x_{n+1} \in \mathbb{X}$, is defined by

$$\bar{\rho}_n(y; x_{n+1}) := \mathbb{E}_n(\rho_{n+1}(y) \mid X_{n+1} = x_{n+1}). \quad (10)$$

(Note that $\bar{\rho}_n(y; x_{n+1}) \neq \rho_n(y)$ because of the implicit dependency of $\rho_{n+1}(y)$ on the future maximum M_{n+1} .)

It turns out that $\bar{\rho}_n(y; x_{n+1})$ can be expressed in closed form, as a function of the posterior mean and covariance of ξ , using the special functions Φ , the cumulative distribution function of the univariate standard normal distribution, and Φ_2 , the cumulative distribution function of the bivariate standard normal distribution. To see this, observe that

$$(\xi(y) - M_{n+1})_+ = \tilde{M}_{n+2} - M_{n+1} = (\tilde{M}_{n+2} - M_n) - (M_{n+1} - M_n), \quad (11)$$

where $\widetilde{M}_{n+2} = \max(M_{n+1}, \xi(y))$. Therefore, we have

$$\bar{\rho}_n(y; x_{n+1}) = \mathbb{E}_n((\xi(y) - M_{n+1})_+ \mid X_{n+1} = x_{n+1}) = \rho_{n,2}(x_{n+1}, y) - \rho_n(x_{n+1}), \quad (12)$$

where $\rho_{n,r}$ denotes the r -point expected improvement criterion:

$$\rho_{n,r}(x_{n+1}, \dots, x_{n+r}) := \mathbb{E}_n(M_{n+r} - M_n \mid X_{n+k} = x_{n+k}, 1 \leq k \leq r). \quad (13)$$

Equation (12) makes it possible to compute $\bar{\rho}_n(y; x_{n+1})$ using the closed-form expression obtained for the multi-point EI by Chevalier and Ginsbourger (2013).

Assuming that $\lambda(\mathbb{X}) < +\infty$, a simple idea for the computation of the integral over \mathbb{X} in (9) is to use a Monte Carlo approximation:

$$\mathfrak{S}_n(x_{n+1}) \approx \frac{\lambda(\mathbb{X})}{m} \sum_{i=1}^m \bar{\rho}_n(Y_i; x_{n+1})$$

where $(Y_i)_{1 \leq i \leq m}$ is a sequence of independent random variables distributed according to $\lambda(\cdot) / \lambda(\mathbb{X})$. Since \mathfrak{S}_n has also to be minimized over \mathbb{X} , we can also use the sample $(Y_i)_{1 \leq i \leq m}$ to carry out a simple stochastic optimization. In practice however, we would recommend to use a more advanced sequential Monte Carlo method, in the spirit of that described in Benassi et al. (2012) and Benassi (2013), to carry out both the integration and the optimization steps.

Remark. Equations (9)–(13) are easily generalized to batch sequential optimization. Define a multi-point EEI by

$$\bar{\rho}_{n,r}(y; x_{n+1}, \dots, x_{n+r}) := \mathbb{E}_n(\rho_{n+r}(y) \mid X_{n+k} = x_{n+k}, 1 \leq k \leq r).$$

We have

$$\bar{\rho}_{n,r}(y; x_{n+1}, \dots, x_{n+r}) = \rho_{n,r+1}(x_{n+1}, \dots, x_{n+r}, y) - \rho_n(x_{n+1}, \dots, x_{n+r}).$$

Then, we can express a multi-point version of the sampling criterion (7) as

$$\mathfrak{S}_{n,r}(x_{n+1}, \dots, x_{n+r}) = \int_{\mathbb{X}} \bar{\rho}_{n,r}(y; x_{n+1}, \dots, x_{n+r}) \lambda(dy).$$

4. Numerical study

The numerical results presented in this section have been obtained with STK (Bect et al., 2014), a free GPL-licensed Matlab/Octave kriging toolbox.

First, we present a simple one-dimensional illustration, whose aim is to contrast qualitatively the behaviour of a sampling strategy based on the EI^2 criterion \mathfrak{S}_n with that of the classical EI-based strategy. Figure 2 depicts a situation where there is a large expected improvement in a small region of the search domain, and a smaller expected improvement over a large region of the search domain. In such a situation, the new sampling criterion \mathfrak{S}_n favors the large region with a smaller expected improvement, thereby inducing a better exploration of the search domain than ρ_n .

Figure 3 represents, for both strategies, the average approximation error obtained on a testbed of 2700 sample paths of a Gaussian process on \mathbb{R}^d , $d = 3$, with zero-mean and isotropic Matérn covariance function, simulated on a set of $m = 1000$ points in $[0, 1]^d$. The isotropic form of the Matérn covariance on \mathbb{R}^d may be written as $k(x, y) = \sigma^2 r_\nu(\|x - y\|/\beta)$, with $r_\nu : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that, $\forall h \geq 0$,

$$r_\nu(h) = \frac{1}{2^{\nu-1}\Gamma(\nu)} (2\nu^{1/2}h)^\nu \mathcal{K}_\nu(2\nu^{1/2}h),$$

where Γ is the Gamma function and \mathcal{K}_ν is the modified Bessel function of the second kind of order ν . Here, $\sigma^2 = 1.0$, $\beta = (4 \cdot 10^{-2} \Gamma(d/2 + 1) / \pi^{d/2})^{1/d} \approx 0.2$ and $\nu = 6.5$. For each optimization strategy, we use the same covariance function for ξ than that used to generate the sample paths in the testbed. Before running the optimization strategies, an initial evaluation point x_1 is set at the center of $[0, 1]^d$. For each sample path f , and each $n \geq 1$, the estimator x_n^* of x^* is defined as $x_n^* = \operatorname{argmax}_{x \in \{x_1, \dots, x_n\}} f(x)$. Thus, $\|x^* - x_n^*\|$ is not a decreasing function of n in general. Figure 3 shows that the approximation errors $M - M_n$ and $\|x^* - x_n^*\|$ decrease approximately at the same rate for both strategies; however, the Euclidean distance of x_n^* to x^* is significantly smaller in the case of the new strategy.




Figure 2: Assessment of the behavior of the sampling criterion \mathfrak{S}_n (bottom, right) against that of ρ_n (bottom, left). The objective is to maximize the function $f : x \in [-1, 1] \mapsto (0.8x - 0.2)^2 + \exp(-\frac{1}{2}|x + 0.1|^{1.95}/0.1^{1.95}) + \exp(-\frac{1}{2}(2x - 0.6)^2/0.1) - 0.02$ (top, dashed line). Evaluations points are represented by squares; the posterior mean $\widehat{\xi}_n$ is represented by a solid line; 95% credible intervals computed using s_n are represented by gray areas. The next evaluation point will be chosen at the minimum of the sampling criterion (vertical solid line).

References

- Bect, J., Ginsbourger, D., Li, L., Picheny, V., Vazquez, E., 2012. Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing* 22 (3), 773–793.
- Bect, J., Vazquez, E., et al., 2014. STK: a Small (Matlab/Octave) Toolbox for Kriging. Release 2.1.
URL <http://kriging.sourceforge.net>
- Benassi, R., 2013. Nouvel algorithme d’optimisation bayésien utilisant une approche monte-carlo séquentielle. Ph.D. thesis, Supélec.
- Benassi, R., Bect, J., Vazquez, E., 2012. Bayesian optimization using sequential Monte Carlo. In: *Learning and Intelligent Optimization*. 6th International Conference, LION 6, Paris, France, January 16-20, 2012, Revised Selected Papers. Vol. 7219 of *Lecture Notes in Computer Science*. Springer, pp. 339–342.
- Chevalier, C., Bect, J., Ginsbourger, D., Vazquez, E., Picheny, V., Richet, Y., 2013. Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics*, 22 pages, accepted for publication, posted online: 21 Nov 2013.
- Chevalier, C., Ginsbourger, D., 2013. Fast computation of the multi-points expected improvement with applications in batch selection. In: Nicosia, G., Pardalos, P. (Eds.), *Learning and Intelligent Optimization*. LNCS. Springer, pp. 59–69.
- Jones, D. R., Schonlau, M., Welch, W. J., 1998. Efficient global optimization of expensive black-box functions. *J. Global Optim.* 13, 455–492.
- Mockus, J., 1989. *Bayesian approach to Global Optimization: Theory and Applications*. Kluwer Acad. Publ., Dordrecht-Boston-London.
- Mockus, J., Tiesis, V., Žilinskas, A., 1978. The application of Bayesian methods for seeking the extremum. In: Dixon, L., Szego, G. (Eds.), *Towards Global Optimization*. Vol. 2. North Holland, New York, pp. 117–129.
- Picheny, V., 2014a. Multiobjective optimization using Gaussian process emulators via stepwise uncertainty reduction. arXiv:1310.0732 (to appear in *Statistics and Computing*).
- Picheny, V., 2014b. A stepwise uncertainty reduction approach to constrained global optimization. In: *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014, Reykjavik, Iceland. Vol. 33. JMLR: W&CP, pp. 787–795.
- Villemonteix, J., Vazquez, E., Walter, E., 2009. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization* 44 (4), 509–534.




Figure 3: Approximation errors of x^* (left) and M (right) using the sampling criteria \mathbf{S}_n (solid line) and ρ_n (dashed line), as a function of the number of evaluations n . More precisely, each plot represents an average approximation error obtained on a testbed of 2700 sample paths of a Gaussian process on \mathbb{R}^3 , with zero-mean and isotropic Matérn covariance function, simulated on a set of 1000 points in $[0, 1]^3$.